

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
15 March 2001 (15.03.2001)


PCT

(10) International Publication Number
WO 01/18654 A1

- (51) International Patent Classification⁷: G06F 13/00
- (21) International Application Number: PCT/IL00/00540
- (22) International Filing Date:
7 September 2000 (07.09.2000)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/152,849 8 September 1999 (08.09.1999) US
60/175,339 10 January 2000 (10.01.2000) US
09/559,352 27 April 2000 (27.04.2000) US
- (71) Applicant (for all designated States except US): MEL-
LANOX TECHNOLOGIES LTD. [IL/IL]; P.O. Box 86,
20692 Yokneam (IL).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): KAGAN, Michael
[IL/IL]; Hashomer Street 71, 30900 Zichron Yaakov (IL).
CRUPNICOFF, Diego [AR/IL]; Sitvanit Street 9, 34793
Haifa (IL). GABBAY, Freddy [IL/IL]; Derech Hashalom
Street 75/2, 67942 Tel Aviv (IL). ROTTENBERG, Shi-
mon [IL/IL]; Dubnov Street 27, 64597 Tel Aviv (IL).
- (74) Agents: COLB, Sanford, T. et al.; Sanford T. Colb & CO.,
P.O. Box 2273, 76122 Rehovot (IL).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ,
DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR,
HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR,
LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ,
NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM,
TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

[Continued on next page]

(54) Title: SYNCHRONIZATION OF INTERRUPTS WITH DATA PACKETS



(57) Abstract: A method and apparatus for conveying data over a packet-switching network (26). Data are received from a peripheral device (25) for transmission via the network to a memory (22) associated with a central processing unit (CPU) (21), followed by an interrupt signal from the peripheral device associated with the data. One or more data packets containing the data are sent over the network to a host network interface (32) serving the memory and the CPU, followed by an interrupt packet sent over the network to the host network interface. Responsive to the interrupt packet, an interrupt input of the CPU is asserted only after the one or more data packets have arrived at the host network interface.



(84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

— *With international search report.*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SYNCHRONIZATION OF INTERRUPTS WITH DATA PACKETS

CROSS-REFERENCE TO RELATED APPLICATION

This application claims the benefit of U.S. Provisional Patent Application 60/152,849, filed September 8, 1999, and of U.S. Provisional Patent Application 60/175,339, filed January 10, 2000. Both of these co-pending applications are assigned to the assignee of the present patent application and are incorporated herein by reference.

FIELD OF THE INVENTION

The present invention relates generally to computing systems, and specifically to systems that use packet-switching fabrics to connect a computer host to peripheral devices.

BACKGROUND OF THE INVENTION

In current-generation computers, the central processing unit (CPU) is connected to the system memory and to peripheral devices by a parallel bus, such as the ubiquitous Peripheral Component Interface (PCI) bus. As data path-widths grow, and clock speeds become faster, however, the parallel bus is becoming too costly and complex to keep up with system demands. In response, the computer industry is moving toward fast, packetized, serial input/output (I/O) bus architectures, in which computing hosts and peripheral are linked by a switching network, commonly referred to as a switching fabric. A number of architectures of this type have been proposed, including "Next Generation I/O" (NGIO) and "Future I/O" (FIO), culminating in the "InfiniBand" architecture, which has been advanced by a consortium led by a group of industry leaders (including Intel, Sun, Hewlett Packard, IBM, Compaq, Dell and Microsoft). Storage Area Networks (SAN) provide a similar, packetized, serial approach to high-speed storage access, which can also be implemented using an InfiniBand fabric.

In a parallel bus-based computer system, when a peripheral device needs to deliver data to the CPU, it typically writes the data to the memory over the bus, using direct memory access. When the peripheral has finished writing, it asserts an interrupt to the CPU on one of the interrupt lines of the bus. Bus arbitration ensures that the CPU will not attempt to read the data from the memory until the writing of the data is complete. On the other hand, when the peripheral device and the CPU are connected by a packet-switching fabric, such as an InfiniBand fabric, they operate asynchronously. Furthermore, the data sent to the memory and the interrupt to the CPU travel over different paths, or channels. Typically, a separate line or

channel is provided to connect the interrupt pin of the peripheral device to an interrupt controller of the CPU, bypassing the switching fabric. Therefore, there is no *a priori* assurance that all of the data will have been written to the memory before the CPU begins reading.

The "race" between the interrupt path and the data path can result in errors (as when a CPU read stalls the data). Care must therefore be taken to synchronize data and interrupt handling and to make sure that the data have been completely written to the memory before the CPU attempts to read it.

A common solution in this situation is to program the CPU to access the peripheral device before accessing the memory, typically by performing a "configuration read" from the peripheral device. In this mode of operation, after the peripheral device has asserted the interrupt to the CPU (indicating that the last item of data has been sent to the memory), the CPU issues a read request through the switching fabric, to read an interrupt cause register in the peripheral device. The peripheral device responds to the read request by sending a packet containing the interrupt cause to the CPU over the same channel as it used to send the data to the memory. Since packets are ordered within a channel, the response to configuration read arrives at the CPU after all of the previous writes have been flushed to memory. The CPU begins to read the data from the memory only after it has received the interrupt cause packet back from the peripheral device. The configuration read thus serves two crucial purposes: it provides the CPU with the cause information that it needs in order to serve the interrupt, and it ensures that the CPU reads the memory only after all of the data have been written there.

This scheme has a number of serious performance drawbacks, however. Every interrupt sent by the peripheral device necessitates an additional exchange of messages through the switching fabric between the CPU and peripheral device. The exchange adds substantial latency – typically 10 microseconds or more - every time the CPU must service an interrupt. Furthermore, since configuration reads are used as synchronization barriers, the CPU is stalled from the moment the configuration read request is issued until its response has arrived. Valuable CPU time is therefore wasted waiting for the interrupt cause to be retrieved.

U.S. Patent 5,689,713, whose disclosure is incorporated herein by reference, describes a method for interrupt request handling in a packet-switched computer system. The system may include a number of interrupt sources, which direct interrupts to any of a number of interrupt handlers. A system controller acts as an intermediary between interrupting devices and

“interruptees.” It includes an interrupt queue coupled to each interrupt source for receiving multiple interrupt requests, and an output queue coupled to each interrupt handler. The controller thus enables asynchronous data from multiple sources to be conveyed across a packet-switched interconnection, while providing a dedicated channel for interrupts associated with the data packets.

SUMMARY OF THE INVENTION

It is an object of the present invention to provide an improved method and system for passing data packets and associated interrupts through a switching fabric.

It is a further object of some aspects of the present invention to provide a method and system for communication between a CPU and peripheral devices via a switching fabric that ensures proper synchronization between data and interrupts transmitted over the fabric.

It is still a further object of some aspects of the present invention to provide a method and system for communication between a CPU and peripheral devices via a switching fabric that reduces latency and processing time required for servicing of interrupts by the CPU.

In preferred embodiments of the present invention, a CPU and a peripheral device are linked to a packet-switching fabric by respective host and target network interfaces. The target interface receives data over a local bus from the peripheral device, for transmission in the form of packets to a system memory associated with the CPU. After sending the data, the peripheral device asserts an interrupt. The interrupt from the device is connected to an interrupt input of the target interface, rather than directly to the CPU or to a central system controller, as in systems known in the art. In response to the interrupt, the target interface reads the interrupt cause from the peripheral device, and then sends a special interrupt packet, including the interrupt cause, to the host interface. Preferably, the target interface sends the interrupt packet on the same channel as it sent the data packets, i.e., over the same “virtual lane,” or route, and with the same priority as the data packets. It thus assures that the host interface will receive the interrupt packet only after it has received all of the preceding data packets.

Upon receiving the interrupt packet, the host interface places the interrupt cause in a predefined register in the memory. An interrupt signal is then sent from the host interface to an interrupt input of the CPU. Upon receiving the signal, the CPU checks to ensure that the host interface has finished writing all of the data from the peripheral device to the memory. This

check serves a similar purpose to the configuration read described in the Background of the Invention. Only after completing the check does the CPU read the interrupt cause and begin processing the data in the memory. The CPU performs all of these steps locally, communicating with the host interface and memory over a local system bus, with latency on the order of nanoseconds, rather than having to exchange messages with the peripheral device through the switching fabric, taking many microseconds. As a result, interrupt response latency is minimized, and the CPU does not waste precious time and resources waiting for the configuration read response.

In preferred embodiments of the present invention, the switching fabric comprises an InfiniBand network, and the host and target interfaces respectively comprise host and target channel adapters. It will be appreciated, however, that the principles of the present invention may similarly be applied to transmission of interrupts through substantially any packet-switched network.

There is therefore provided, in accordance with a preferred embodiment of the present invention, a method for conveying data over a packet-switching network, including:

receiving data from a peripheral device for transmission via the network to a memory associated with a central processing unit (CPU);

receiving an interrupt signal from the peripheral device associated with the data;

sending one or more data packets containing the data over the network to a host network interface serving the memory and the CPU; and

sending an interrupt packet over the network to the host network interface, responsive to which an interrupt input of the CPU is asserted only after the one or more data packets have arrived at the host network interface.

Typically, receiving the data includes receiving parallel data over a local bus from the peripheral device. Additionally or alternatively, receiving the data includes receiving data to be written to the memory by direct memory access.

Preferably, sending the interrupt packet includes reading a cause of the interrupt from the peripheral device, and incorporating the cause in the interrupt packet. Further preferably, the method includes receiving the interrupt packet at the host network interface, and writing the cause to a predetermined address in the memory, to be read by the CPU after the interrupt input is asserted.

In a preferred embodiment, sending the interrupt packet includes sending the interrupt packet after receiving an acknowledgment from the memory that the data have been written thereto.

Preferably, sending the one or more data packets includes sending the data packets over
5 a selected channel through the network, and sending the interrupt packet includes sending the interrupt packet over the selected channel following the data packets.

Further preferably, the method includes:

receiving the data packets and the interrupt packet at the host network interface;

conveying the data in the packets for delivery to the memory over a local bus coupling
10 the host network interface to the memory and the CPU; and

notifying the CPU when all of the data have been conveyed.

Most preferably, conveying the data in the packets includes passing the data to a system controller on the bus, and notifying the CPU includes informing the CPU when an acknowledgment is received by the host network interface from the system controller, typically
15 by asserting the interrupt input of the CPU after the acknowledgment from the system controller has been received. Additionally or alternatively, notifying the CPU includes asserting the interrupt input of the CPU responsive to receiving the interrupt packet at the host network interface.

There is also provided, in accordance with a preferred embodiment of the present
20 invention, network interface apparatus, including:

a target channel adapter, which is operative to receive data from a peripheral device for transmission via a packet-switching network to a memory associated with a central processing unit (CPU) and to send one or more data packets containing the data over the network to a host network interface serving the memory and the CPU; and

25 a target interface processor, adapted to receive an interrupt signal from the peripheral device associated with the data, and to send an interrupt packet over the network to the host network interface, responsive to which an interrupt input of the CPU is asserted only after the one or more data packets have arrived at the host network interface.

There is further provided, in accordance with a preferred embodiment of the present
30 invention, network interface apparatus, including:

Although for simplicity, only a single interrupt line from unit 28 to controller 38 is shown in Fig. 1, the unit preferably comprises multiple interrupt lines. These lines can be actuated selectively by controller 36 so as to send multiple, different interrupts to CPU 21 depending on the content of interrupt packets received by the HCA. Alternatively or
5 additionally, the different interrupt lines may be used to signal other host devices that are linked to bus 50.

Peripheral device 25 is coupled to fabric 26 by a target network interface unit 40, similar in structure to unit 28. A target channel adapter (TCA) 42 in unit 40 interfaces via an I/O bus 53 with device 25. Typically, although not necessarily, bus 53 comprises a PCI bus,
10 like bus 50. A switch 44 links the TCA to the switching fabric. A target unit controller 46, similar to controller 36, acts as FSA to TCA 42 and switch 44 and also has a suitable input to receive signals from interrupt output 48 of device 25.

Fig. 2 is a flow chart that schematically illustrates a method by which target interface unit 40 processes and transmits data from peripheral device 25 to HCA 32 over fabric 26, in
15 accordance with a preferred embodiment of the present invention. At a data writing step 60, device 25 writes data via bus 53 to TCA 42, to be conveyed by direct memory access to memory 22. The peripheral device assigns a priority to the data to be transmitted and informs the TCA of this priority. At a data sending step 62, the TCA packetizes the data and sends it over fabric 26 to the address of HCA 32, with the priority assigned by the peripheral device. A
20 packet header instructs the HCA to write the data to memory 22. Preferably, the TCA negotiates with switch 44 and fabric 26 to assign a fixed route for all of the packets through the fabric. Such a route, together with the priority of the packets, is referred to herein as a channel. InfiniBand specifies that packets travelling over the same channel are always kept in their original order.

25 When device 25 has finished posting to TCA 42 all of the data that it has to send, it asserts interrupt output 48, at an interrupt assertion step 64. At the same time, the peripheral device places the cause for the interrupt (in this case, to instruct CPU 21 to read the data from memory 22) in an interrupt cause register 49. In systems known in the art, when the CPU receives the interrupt, it must communicate with the peripheral device in order to read this
30 register. In system 20, however, the interrupt signal is received by controller 46, which instructs TCA 42 to read the interrupt cause from register 49, at a cause reading step 66.

Based on the interrupt cause information read by the TCA, controller 46 constructs an interrupt packet containing the interrupt cause information, at an interrupt packet sending step 68. The interrupt packet is a management packet addressed to the LID of HCA 32. It is preferably sent by controller 46 over the same channel, or virtual lane, as the data packets, after the last of the data packets has been sent. The interrupt packet also identifies the data with which the interrupt is associated. As a result, when the interrupt packet arrives at its destination, controller 36 will be able to generate an interrupt to CPU 21 that is associated with the appropriate memory write, as described below. Controller 46 assures that interrupt packet is sent to the fabric after all of the data packets have already been accepted for sending. It thus ensures that HCA 32 will receive the interrupt packet only after it has received all of the data packets.

As an alternative, controller 46 may delay sending the interrupt packet until TCA 42 receives an acknowledgment from memory 22 that it has received all of the data. This approach introduces additional delay before CPU 21 can receive and act upon the interrupt, but it obviates the need to ensure that the interrupt packet is routed over the same channel as the data packets. Such an approach may be called for in particular when switching fabric 26 comprises a network in which consistent routing and ordering are not necessarily maintained among successive packets. This approach can also be used when the interrupt path and data path are not the same, and fork at an earlier stage than in Fig. 1. Such path incongruity may occur, for example, when the device writing data to the memory is different from the device asserting the interrupt to the CPU. Sometimes it is also desirable to send interrupts on different (high-priority) routes, because data routes can be congested, causing interrupt messages to get stuck behind data.

Fig. 3 is a flow chart that schematically illustrates a method by which data and accompanying interrupt packets are received and processed by host interface unit 28 and CPU 21, in accordance with a preferred embodiment of the present invention. At a packet reception step 70, HCA 32 receives the data and interrupt packets sent from target interface unit 40. The HCA posts the data in the data packets via bus 50 to a buffer 58 of system controller 24. The system controller proceeds to write the data from its buffer to the appropriate addresses in memory 22, as is known in the art. The HCA passes the interrupt packet to controller 36 for

decoding, at an interrupt processing step 72. The controller extracts the cause of the interrupt and posts this information, via HCA 32, to an interrupt cause register 56 in memory 22.

Before CPU 21 services the interrupt represented by the interrupt packet, it is necessary to ensure that all of the associated data have been written to memory 22, at a delivery completion step 74. In the case that controller 46 of target interface unit 40 is programmed to send the interrupt packet only after receiving the acknowledgment from memory 22, as described above, this problem is already solved. Otherwise, controller 36 preferably waits to assert the interrupt until system controller 24 has acknowledged to HCA 32 that it has received all of the data. In response to this acknowledgment, controller 36 sends an interrupt signal to interrupt controller 38, at an interrupt assertion step 76. The interrupt controller actuates interrupt input 27 of CPU 21, to inform the CPU that an interrupt has arrived from HCA 32. In response to the interrupt, the CPU preferably sends a dummy read command to the HCA, in order to ensure that buffer 58 is flushed to memory 22 before the CPU itself begins to process the data in the memory.

As a further alternative, as long as it is assured that the interrupt packet reached HCA 32 after the last of the data packets (which will be the case when all of the packets are sent over the same channel, as described above), controller 36 may send the interrupt signal to interrupt controller 38 immediately, without waiting for an acknowledgment from system controller 24. In this case, upon receiving the interrupt, CPU 21 preferably sends a "fence" command to HCA 32. This command instructs the HCA to mark the last packet currently in its receive queue, and to inform the CPU when this last packet has been written to system controller 24. At this point, the CPU can send its dummy read command and begin processing the data in the memory.

Once it is assured that all of the relevant data have reached their destination in memory 22, CPU 21 reads the cause of the current interrupt from register 56, at a cause reading step 78. Based on this information, the CPU processes the data that peripheral device 25 has placed in the memory, at a data processing step 80. Unlike methods of interrupt processing known in the art, all of the steps in the method of Fig. 3 are carried out locally, typically over busses 50 and 52, without the need for messages to traverse fabric 26.

It will be appreciated that the preferred embodiments described above are cited by way of example, and that the present invention is not limited to what has been particularly shown

and described hereinabove. Rather, the scope of the present invention includes both combinations and subcombinations of the various features described hereinabove, as well as variations and modifications thereof which would occur to persons skilled in the art upon reading the foregoing description and which are not disclosed in the prior art.

CLAIMS

1. A method for conveying data over a packet-switching network, comprising:
receiving data from a peripheral device for transmission via the network to a memory
associated with a central processing unit (CPU);
5 receiving an interrupt signal from the peripheral device associated with the data;
sending one or more data packets containing the data over the network to a host
network interface serving the memory and the CPU; and
sending an interrupt packet over the network to the host network interface, responsive
to which an interrupt input of the CPU is asserted only after the one or more data packets have
10 arrived at the host network interface.
2. A method according to claim 1, wherein receiving the data comprises receiving parallel
data over a local bus from the peripheral device.
3. A method according to claim 1, wherein receiving the data comprises receiving data to
be written to the memory by direct memory access.
- 15 4. A method according to claim 1, wherein sending the interrupt packet comprises reading
a cause of the interrupt from the peripheral device, and incorporating the cause in the interrupt
packet.
5. A method according to claim 4, and comprising receiving the interrupt packet at the
host network interface, and writing the cause to a predetermined address in the memory, to be
20 read by the CPU after the interrupt input is asserted.
6. A method according to any of the preceding claims, wherein sending the interrupt
packet comprises sending the interrupt packet after receiving an acknowledgment from the
memory that the data have been written thereto.
7. A method according to any of claims 1-5, wherein sending the one or more data packets
25 comprises sending the data packets over a selected channel through the network, and wherein
sending the interrupt packet comprises sending the interrupt packet over the selected channel
following the data packets.
8. A method according to any of claims 1-5, and comprising:
receiving the data packets and the interrupt packet at the host network interface;

conveying the data in the packets for delivery to the memory over a local bus coupling the host network interface to the memory and the CPU; and
notifying the CPU when all of the data have been conveyed.

9. A method according to claim 8, wherein conveying the data in the packets comprises
5 passing the data to a system controller on the bus, and wherein notifying the CPU comprises informing the CPU when an acknowledgment is received by the host network interface from the system controller.

10. A method according to claim 9, wherein informing the CPU comprises asserting the interrupt input of the CPU after the acknowledgment from the system controller has been
10 received.

11. A method according to claim 8, wherein notifying the CPU comprises asserting the interrupt input of the CPU responsive to receiving the interrupt packet at the host network interface.

12. Network interface apparatus, comprising:
15 a target channel adapter, which is operative to receive data from a peripheral device for transmission via a packet-switching network to a memory associated with a central processing unit (CPU) and to send one or more data packets containing the data over the network to a host network interface serving the memory and the CPU; and
a target interface processor, adapted to receive an interrupt signal from the peripheral
20 device associated with the data, and to send an interrupt packet over the network to the host network interface, responsive to which an interrupt input of the CPU is asserted only after the one or more data packets have arrived at the host network interface.

13. Apparatus according to claim 12, wherein the target channel adapter comprises an interface to a local parallel bus linked to the peripheral device, over which the device sends the
25 data.

14. Apparatus according to claim 12, wherein the target channel adapter is operative to read a cause of the interrupt from the peripheral device, and wherein the processor is adapted to incorporate the cause in the interrupt packet.

15. Apparatus according to claim 14, and comprising a host channel adapter, coupled to receive the interrupt packet at the host network interface, and to write the cause to a predetermined address in the memory, to be read by the CPU after the interrupt input is asserted.

16. Apparatus according to any of claims 12-15, wherein the processor is adapted to send the interrupt packet after receiving an acknowledgment from the memory that the data have been written thereto.

17. Apparatus according to any of claims 12-15, wherein the target channel adapter is coupled to send the data packets over a selected channel through the network, and wherein the processor is adapted to send the interrupt packet over the selected channel following the data packets.

18. Apparatus according to claim 17, and comprising a switch coupling the target channel adapter and the processor to the network, wherein the switch comprises a receive queue into which the target channel adapter places the data packets, and wherein the processor is adapted to place the interrupt packet into the receive queue following the data packets.

19. Apparatus according to any of claims 12-15, and comprising a host interface unit, which is coupled to receive the data and interrupt packets transmitted over the network, and is operative to convey the data in the packets for delivery to the memory over a local bus coupled to the memory and the CPU and to notify the CPU when all of the data have been conveyed.

20. Apparatus according to claim 19, wherein the host interface unit is coupled to assert the interrupt to the CPU responsive to the interrupt packet.

21. Apparatus according to any of claims 12-15, wherein the target channel adapter comprises an InfiniBand adapter.

22. Network interface apparatus, comprising:

a host channel adapter, which is operative to receive data packets transmitted over a packet-switching network from a peripheral device, and to convey data from the packets for delivery to a memory associated with a CPU over a local bus that is coupled to the memory and the CPU, and further to receive an interrupt packet sent over the network responsive to an interrupt signal asserted by the peripheral device after sending the data to the network; and

a host interface processor, adapted, responsive to the interrupt packet, to notify the CPU when all of the data have been conveyed to the local bus.

23. Apparatus according to claim 22, wherein the host channel adapter is operative to convey the data to the memory by direct memory access.

5 24. Apparatus according to claim 22, wherein the host channel adapter is operative to convey the data to a system controller on the bus, and wherein the CPU is notified when an acknowledgment is received by the host channel adapter from the system controller.

10 25. Apparatus according to claim 24, wherein the host interface processor is coupled to assert the interrupt input of the CPU after the acknowledgment from the system controller has been received.

26. Apparatus according to any of claims 22-25, wherein the host interface processor is coupled to assert the interrupt input of the CPU responsive to receipt of the interrupt packet at the host network interface.

15 27. Apparatus according to any of claims 22-25, wherein the host channel adapter comprises an InfiniBand adapter.

1/2




FIG. 1

2/2

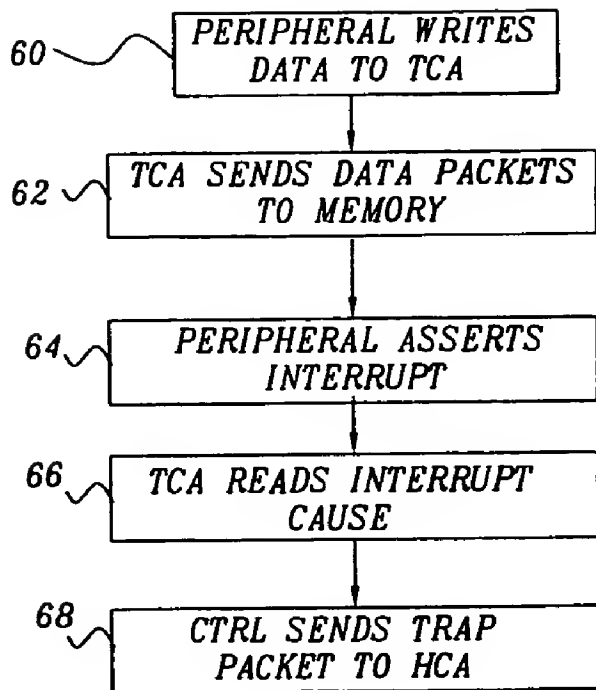
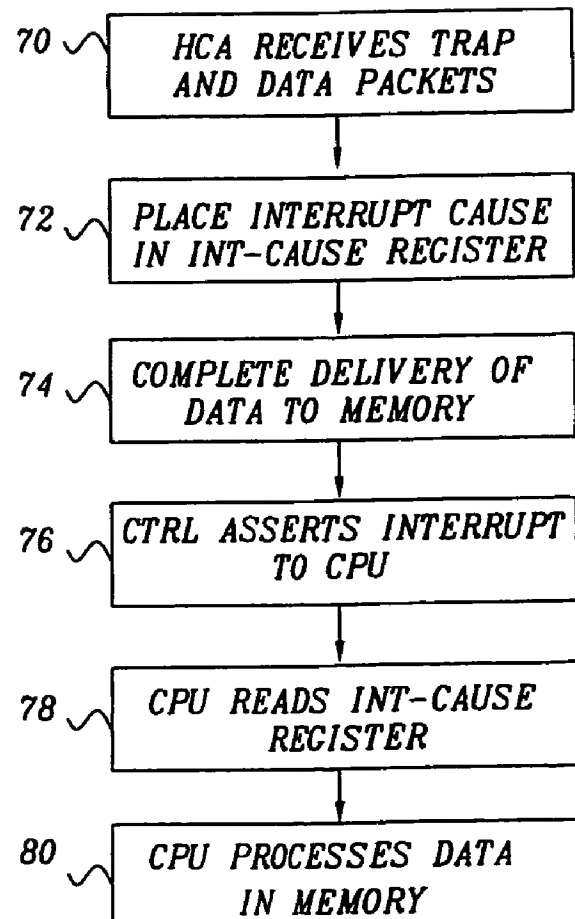


FIG. 2

FIG. 3



INTERNATIONAL SEARCH REPORT

International application No.

PCT/IL00/00540

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 13/00

US CL : 710/260

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 710/260,263,266; 370/426; 709/238,250

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
Please See Continuation Sheet

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,659,758 A (GENTRY et al) 19 August 1997, abstract, figs. 3 and 6, and columns 2, 4-5.	1-3, 7, 8, 11-13, 17, 19, 20, 22-26
A,P	US 6,038,629 A (OGILVIE et al) 14 March 2000, abstract.	1-27
A	US 5,440,545 A (BUCHHOLZ et al) 08 August 1995, abstract and columns 1-2.	1-27
A	US 5,689,713 A (NORMOYLE et al) 18 November 1997, abstract.	1-27



Further documents are listed in the continuation of Box C.



See patent family annex.

Special categories of cited documents:	
* "A" document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E" earlier application or patent published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

15 November 2000 (15.11.2000)

Date of mailing of the international search report

28 DEC 2000

Name and mailing address of the ISA/US

Commissioner of Patents and Trademarks

Box PCT

Washington, D.C. 20231

Facsimile No. (703)305-3230

Authorized officer

Glenn A. Auye James R. Hattchew

Telephone No. (703) 305-3900

INTERNATIONAL SEARCH REPORT

International application No.

PCT/IL00/00540

Continuation of B. FIELDS SEARCHED Item 3: EAST text search: channel adj adapter; packet adj switch\$1; adj network;
interrupt adj2 packet\$1; infiniband
WWW search using Google search engine: infiniband

This Page Blank (uspio)